

07-31-00

A

07/28/00




jc866 U.S. PTO

# UTILITY PATENT APPLICATION AND FEE TRANSMITTAL

(Only for new nonprovisional applications under 37 C.F.R. 1.53(b))

Attorney Docket No.	12780-1015	Total Pages	
First Named Inventor or Application Identifier			
Ian R. Finlay			
Express Mail Label No.	EE087691528US		

jc857 U.S. PTO  
09/28/00

<b>APPLICATION ELEMENTS</b> See MPEP chapter 600 concerning utility patent application contents.		<b>ADDRESS TO:</b> Assistant Commissioner for Patents Box Patent Application Washington, D.C. 20231	
1. <input checked="" type="checkbox"/> Fee Transmittal Form (Submit an original, and a duplicate for fee processing.)		6. <input type="checkbox"/> Microfiche Computer Program (Appendix)	
2. <input checked="" type="checkbox"/> Specification Total Pages <input type="text" value="15"/> - Descriptive title of the invention - Background of the Invention - Brief Summary of the Invention - Brief Description of the Drawings (if filed) - Detailed Description - Claim(s) - Abstract of the Disclosure		7. <input type="checkbox"/> Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary) a. <input type="checkbox"/> Computer readable copy b. <input type="checkbox"/> Paper copy (identical to computer copy) c. <input type="checkbox"/> Statement verifying identity of above copies	
3. <input checked="" type="checkbox"/> Drawing(s) (35 U.S.C. 113) Total Sheets <input type="text" value="4"/>		8. <input type="checkbox"/> Assignment Papers (cover sheet & documents)	
4. <input checked="" type="checkbox"/> Oath or Declaration Total Pages <input type="text" value="2"/> a. <input checked="" type="checkbox"/> (Unsigned) copy b. <input type="checkbox"/> Copy from a prior application (37 C.F.R. 1.63(d)) (for continuation/divisional with Box 17 completed) (Note Box 5 below) i. <input type="checkbox"/> Deletion of Inventor(s) Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. 1.63(d)(2) and 1.33(b).		9. <input type="checkbox"/> 37 C.F.R. 3.73(b) Statement <input type="checkbox"/> Power of Attorney 10. <input type="checkbox"/> English Translation Document (if applicable) 11. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input type="checkbox"/> Copies of IDS Citations 12. <input type="checkbox"/> Preliminary Amendment	
5. <input type="checkbox"/> Incorporation By Reference (useable if Box 4b is checked). The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.		13. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) (Should be specifically itemized) 14. <input type="checkbox"/> Small Entity Statement(s) <input type="checkbox"/> Statement filed in prior appn. Status still proper and desired. (unsigned) 15. <input checked="" type="checkbox"/> Certified Copy of Priority Document(s) (If foreign priority is claimed)	
17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information <input type="checkbox"/> Continuation <input type="checkbox"/> Divisional <input type="checkbox"/> Continuation-in-part (CIP) of prior application No. : (fill in)			
<b>18. CORRESPONDENCE ADDRESS</b> <input type="checkbox"/> Customer Number or Bar Code Label <input checked="" type="checkbox"/> New correspondence address below			

OPPENHEIMER WOLFF & DONNELLY LLP  
 2029 Century Park East, Suite 3800  
 Los Angeles, California 90067 U.S.A.  
 Telephone (310) 788-5000 • Fax (310) 788-5100  
 Attn: Spyros James Lazaris

## DIRECT CALL THREADED CODE

By

Ian R. Finlay

Douglas James Doole

5

BACKGROUND OF THE INVENTIONField of the Invention

The present invention relates to relational database management systems, and more particularly to a method for optimizing execution of queries in a relational database management system.

10 Description of the Related Art

A database management system (DBMS) comprises the combination of an appropriate computer, direct access storage devices (DASD) or disk drives, and database management software. A relational database management system is a DBMS which uses relational techniques for storing and retrieving information. The relational database management system or RDBMS, such as the DB2 product from IBM, comprises computerized information storage and retrieval systems in which data is stored on disk drives or DASD for semi-permanent storage. The data is stored in the form of tables which comprise rows and columns. Each row or tuple has one or more columns.

The RDBMS is designed to accept commands to store, retrieve, and delete data. One widely used and well-known set of commands is based on the Structured Query Language or SQL. The term query refers to a set of commands in SQL for retrieving data from the RDBMS. The definitions of SQL provide that a RDBMS should respond to a particular query with a particular set of data given a specified database content. SQL however does not specify the actual method to find the requested information in the tables on the disk drives. There are many ways in which a query can be processed and each consumes a different amount of processor and

input/output access time. The method in which the query is processed, i.e. query execution plan, affects the overall time for retrieving the data. The time taken to retrieve data can be critical to the operation of the database. It is therefore important to select a method for finding the data requested in a RDBMS which minimizes the computer and disk access time, thereby optimizing the cost of doing the query.

To execute a SQL statement, such as "select \* from t1", in a RDBMS program like DB2 the statement is presented to the SQL optimizer. The SQL optimizer parses, tokenizes and semantically analyzes the statement, transforming it into the Query Graph Model (QGM) representation of the statement. The QGM representation is then processed to perform a number of heuristic optimizations. The output of this pass is then fed to the cost-based planning stage. The cost-based planning stage processes the optimized QGM, producing an access plan, based on LOW LEVEL Plan Operators (LOLEPOPs). The plan produced by the cost-based optimizer is then processed by a Code Generation module (CODGEN) to produce an OPERATION CODE (OPCODE) based access plan, which can be processed at runtime by a Relational Database System (RDS).

In prior versions of RDBMS programs, such as DB2 (Versions 5.2 and older), the OPCODE based access plan is interpreted at runtime by the Relational Database System (RDS). The Relational Database System examines each OPCODE, and looks up the function which is called to process the OPCODE and its operands. The processing for the OPCODE includes loading the OPCODE's operands and making decisions based on information associated with the OPCODE that was provided at CODGEN time. These decisions are made repeatedly, each time the OPCODE processing function is called, and direct the function of the OPCODE. An alternative implementation involved producing multiple OPCODEs for these similar functions. This approach still results in considerable duplication in underlying OPCODE processing.

It will be appreciated that one of the principle problems with existing RDBMS programs, such as the DB2 product, is the fact that the RDBMS includes an interpreter which executes during runtime. Since the interpreter translates and runs the OPCODEs at the same time, operation during runtime is considerably slower than for a compiler based implementation. In view of the costs associate with replacing existing interpreter-based RDBMS programs, there remains a need for a mechanism which can improve the slower runtime performance of the interpreter phase in such systems.

SUMMARY OF THE INVENTION

The pre-pass mechanism of the present invention replaces the repeated looking up of the function to call to process the OPCODE and the function's operands, and any decisions that need to be made repeatedly (i.e. static decisions), during the interpreter phase of execution. The pre-pass mechanism comprises a pre-processing function which replaces or augments the OPCODE, and any static decisions, with a pointer to the function to call to perform the operation specified by the OPCODE, or a pointer to an intermediate function with an auxiliary data structure, or a pointer to an auxiliary data structure, wherein the auxiliary data structure includes a pointer to the function to call to perform the operation specified by the OPCODE. Advantageously, the pointers are called without additional lookup. The intermediate function to call to perform the function specified by the OPCODE may include processing operations and static decision making.

In a first aspect, the present invention provides a method for pre-processing an access plan generated for a query in a relational database management system. The access plan has a plurality of operation codes including a first operation code and is stored in memory, each of the operation codes being associated with one or more executable functions for performing the query. The method comprises the steps of: (a) retrieving the access plan from memory; (b) determining an executable function associated with the first operation code in the access plan; (c) augmenting the first operation code in the access plan with a pointer to the executable function; (d) repeating steps (b) and (c) for the remaining operation codes in the access plan; (e) storing the modified access plan in memory.

In a second aspect, the present invention provides a relational database management system for use with a computer system wherein queries are entered by a user for retrieving data from tables, the relational database management system includes a query optimizer for generating an access plan associated with the queries entered by the user, the relational database management system comprises: (a) means for retrieving the access plan generated for a query from memory; (b) means for determining an executable function associated with each of the operation codes in the access plan; (c) means for augmenting the operation codes in the access plan with a pointer to the associated executable function; (d) means for storing a modified access plan in memory.

In another aspect, the present invention provides a computer program product for use on a computer wherein queries are entered by a user for retrieving data in a relational database management system having a query optimizer for generating an access plan for executing the query, the computer program product comprises: a recording medium; means recorded on the medium for instructing the computer to perform the steps of, (a) retrieving the access plan from memory; (b) determining an executable function associated with the first operation code in the access plan; (c) augmenting the operation code in the access plan with a pointer to the executable function; (d) repeating steps (b) and (c) for the remaining operation codes in the access plan; (e) storing the modified access plan in memory.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows in diagram form a computer system for the present invention;

FIG. 2 shows in flowchart form the steps for the interpretation and execution of SQL statements in the context of the present invention;

FIG. 3 shows in flowchart form the steps for translating OPCODE to function pointers according to the present invention;

FIG. 4 shows in flowchart form the steps for execution of the translated OPCODE according to the present invention.

#### DETAILED DESCRIPTION OF THE INVENTION

Reference is first made to FIG. 1 which shows an exemplary computer system 10 which is suitable for use with the present invention. As shown in FIG. 1, the computer system 10 comprises one or more processors connected to one or more electronic storage devices 12 and 14, such as disk drives, that store one or more relational databases. Users of the computer system 10 use a standard operator interface 16, such as OS/2 or other similar interface, to input commands for performing various search and retrieval functions, termed queries, against the databases. In the context of the present invention, these queries conform to the Structured Query Language (SQL) standard, and invoke functions performed by a Relational Database Management System (RDBMS). In the preferred embodiment of the present invention, the

RDBMS software comprises the DB2 product offered by IBM for the MVS or OS/2 operating systems. Those skilled in the art will however recognize that the present invention has application to any RDBMS program that utilizes SQL.

As shown in FIG. 1, the DB2 architecture for the MVS operating system includes three major components: an IMS Resource Lock Manager (IRLM) 18, a System Services module 20, and a Database Services module 22. The IRLM 18 handles locking services because DB2 treats data as a shared resource, thereby allowing any number of users to access the same data simultaneously. As a consequence, concurrency control is required to isolate users and to maintain data integrity. The System Services module 20 controls the overall DB2 program execution environment, which includes managing log data sets 14, gathering statistics, handling startup and shutdown, and providing management support.

At the center of the DB2 architecture is the Database Services module 22. The Database Services module 22 contains several sub-modules, including the Relational Database System (RDS) 24, the Data Manager 26, and other components such as a SQL compiler. These sub-modules support the functions of the SQL language, i.e. definition, access control, retrieval, and update of user and system data.

Reference is next made to FIG. 2 which shows in flowchart form the steps for the interpretation and execution of SQL statements in the context of the present invention. Block 102 represents the input by the user of SQL statements into the computer system 10. Block 103 represents the step of checking to determine if the RDBMS has previously processed the SQL statements input by the user. If the SQL statements have already been processed, then the access path plan is retrieved in Block 107 and processed as described below. If SQL statements have not been processed, then processing is performed starting at Block 104. Block 104 represents the step of compiling or interpreting the SQL statements. Block 106 represents the step of generating a compiled set of runtime structures called an execution or access path plan from the compiled SQL statements. Generally, the SQL statements received as inputs from the user specify only the data that the user wants, but not how to get it. This step considers both the available access paths (i.e. indexes, sequential reads, etc.) and system held statistics on the data to be accessed (i.e. the size of the table, the number of distinct values in a particular column, etc.) to choose what it considers to be the most efficient access path for the query. Block 109 represents the step performed in the Access Plan Manager (included in the RDBMS 24 - FIG. 1)

for storing the access path plan (as generated through Block 106) into a system table which allows future re-use of the access path plan. If the SQL statements (i.e. query) was previously processed (as determined in block 103), then the access path plan is retrieved from the system tables as indicated by Block 107. Block 108 represents the step of determining if the access path plan is still valid. In Block 108, the Access Plan Manager examines the retrieved access path plan to ensure that the plan is still valid, i.e. all indexes, tables, etc. still exist in the RDBMS. If the access path plan is not valid, then a new access path plan needs to be generated starting at Block 104 (as described above). If the access path plan is valid, then the next step in Block 109 involves storing the access path plan in the system tables. The access path plan is also stored in a memory cache as indicated in Block 109. Block 110 represents the execution of the application plan. A pre-pass function according to present invention within block 110 improves runtime performance as will be described below. Block 112 represents the output of the results of the application or access path plan to the user.

In prior art systems as described above, the application or access plan is generated by a cost-based optimizer and then processed by a Code Generation module (included in other components 22 shown in FIG. 1) to produce an Operation Code (OPCODE) access plan. The OPCODE access plan is then stored in the main memory cache for the computer system 10. To execute the SQL statement (block 110), i.e. the access plan, according to the prior art, the computer system 10 would first retrieve the section of code from the memory cache. An instruction pointer (IP) is initialized to point to the first entry in the list of operations, or OPCODEs, that describe what is to be done to execute the SQL statement. The Relational Database System (RDS) uses the OPCODE to look up a corresponding function (F) in a look-up table. The look-up table contains the mapping of every OPCODE supported by the database engine 26 to the executable function that performs the operation associated with the OPCODE. Next the function (F) is called by the RDS to apply the function as required. A "return" code is also returned along with a next instruction pointer (IPNxt). The return code indicates whether execution is to stop or continue, and the next instruction pointer (IPNxt) points to the next operation that is to be executed now that the current operation is complete. The return code typically indicates STOP if the SQL request had been completed. It will be appreciated that with this prior art approach, the OPCODE may be looked up and executed multiple times in the

course of processing an SQL request. Repeated execution is further compounded by nested SQL requests.

According to the present invention, a pre-pass function or mechanism is included in block 110. The pre-pass function pre-processes the access path plan for the SQL statement before  
5 storage in the memory cache. The pre-processing comprises looking up the OPCODEs to determine the corresponding function to execute with each OPCODE. The OPCODE is replaced or augmented with a pointer to the function to call to perform the operation specified by the OPCODE. The pre-processing further includes an examination of the function and runtime environment to determine more precisely which function best suits the operation being requested  
10 for the SQL statement.

According to another aspect of the pre-pass mechanism, an intermediate processing function may be inserted either during this pre-pass stage, or at a later processing stage, to provide additional, auxiliary processing capabilities. The pointer to the function to call to process the operation specified by the OPCODE is replaced by a pointer to the intermediate  
15 processing function and an auxiliary data structure. The auxiliary data structure includes a pointer to store the pointer to the function to call to process the operation specified by the OPCODE. The intermediate function provides the capability to perform additional, auxiliary processing such as gathering statistics on the function which process the operation specified by the OPCODE, or requesting input from a user via the terminal interface 16, before or after  
20 calling the function to process the operation specified by the OPCODE. According to another aspect, the pointer to the function to call to process the operation specified by the OPCODE is augmented with another pointer to the auxiliary data structure.

Reference is made to FIG. 3 which shows in flowchart form the steps for performing a pre-pass function according to the present invention indicated generally by reference 200. The  
25 pre-pass function is performed by the Access Plan Manager (APM) in the RDBMS 24 (FIG. 1) prior to storing the access path plan into the memory cache, but after storing the unaltered access path plan in the system tables. This allows for future reuse of the access path plan should the user input the same SQL statements at a later time. Block 202 represents the first step in the pre-pass function and involves loading the code section from the memory cache. Block 204  
30 represents the step of initializing a loop counter LC to 1 for looping through all of the OPCODEs in the code section. Block 206 represents the steps of looking up the function F corresponding to



the OPCODE being referenced by the loop counter LC. In addition, the function is examined to determine if a more specific function can be used to process this request. If a more specific function exists, it is selected to replace the function F. Block 208 represents the step of storing a call pointer to the function F for the OPCODE being referenced by the loop counter LC.

5 Alternatively, block 208 may also include steps for storing a call pointer to an intermediate processing function, an auxiliary data structure, wherein the auxiliary data structure includes a call pointer to the function F for the OPCODE being referenced by the loop counter LC. The call pointer is stored in a function pointer OP\_F(LC) and is indexed by the current value of the loop counter LC which corresponds to the OPCODE being currently processed. Block 210  
10 represents a decision step which checks the loop counter LC to determine if all the OPCODEs in the code section have been processed. If all the OPCODEs have not been processed, then the loop counter LC is incremented in block 212 and the steps in blocks 206 to 210 are repeated for the remaining OPCODEs. Once all the OPCODEs are processed, the processed code section is stored in the memory cache (Block 214) and the pre-pass function 200 ends (Block 216). The  
15 processed code section includes the function pointers corresponding to the OPCODEs previously stored in the code section, and the code section is ready for execution.

Reference is next made to FIG. 4, which shows the sequence of steps for executing a SQL request which has been processed by the pre-pass function 200 (FIG. 3). As shown in FIG. 4, Block 302 represents the first step which involves retrieving the processed code section from  
20 the memory cache. Block 304 represents the step of initializing an instruction pointer IP to point to the first function entry in the code section. Block 306 represents the step of the RDS calling the function that is being referenced by the function pointer that is pointed to by the instruction pointer OP\_F(IP) to execute the function being referenced. In addition, a return code RC and a next instruction pointer IPNxt are returned in Block 306. The return code RC indicates whether  
25 to stop or continue, and the next instruction pointer IPNxt indicates the next function pointer that is to be called now that the current function has been executed. Block 308 represents the step of checking if the return code RC indicates STOP. The return code RC indicates STOP if the SQL request has been completed. If not completed, the instruction pointer IP is set to the value of the next instruction pointer IPNxt and the steps in Blocks 306 to 308 are repeated.

30 It will be appreciated that the pre-pass mechanism according to the present invention replaces the need for the repeated looking up of the function to call in order to process the

OPCODE and its operands. The pre-pass mechanism also replaces any decisions that need to be made repeatedly (i.e. static decisions). The look up operation to the function is replaced with a pointer to the function, whereas static decisions are replaced by pre-processing operation(s).

There are a number of processing options that can be performed with the pre-pass mechanism. For example, the RDS in the current version of DB2 has a save/restore that is always performed for all new subroutines that start up. These operations are not always required, as the information they save is not always affected by the operations that are performed within a section. The save/restore operations can be moved into entry/exit routines, that are then tacked on to the start and end of the subroutine, thus saving considerable overhead for sections that are repeatedly called, and that do not require the save/restore. As described above, static decisions may also be moved into a pre-processing function, allowing the removal of 'if-then-else' decisions from the normal runtime path. This is particularly beneficial for decisions that are included due to the general nature of some operands, but that rarely are used. Removing such decisions from being made repeatedly offers significant performance improvements for such operations. An example of such a decision is code page translation for string operations. Although operands are often in the same code page, it is still necessary to confirm this, but it need only be confirmed once. If the operands are in the same code page, then a more optimal string operator may be called, otherwise, a more general string operator, which takes into account dissimilar code pages, must be called. Similar conditions exist for predicates, and mathematical operators, where optimal functions may be called to handle like-to-like operands, in place of the general operators which handle dissimilar operands.

The present invention may be embodied in other specific forms without departing from the spirit or essential characteristics thereof. Therefore, the presently discussed embodiments are considered to be illustrative and not restrictive, the scope of the invention being indicated by the appended claims rather than the foregoing description, and all changes which come within the meaning and range of equivalency of the claims are therefore intended to be embraced therein.

WHAT IS CLAIMED IS:

1. A method for pre-processing an access plan generated for a query in a relational database management system, said access plan including a plurality of operation codes, each of said operation codes being associated with one or more executable functions for performing the query, said method comprising the steps of:

(a) determining from the access plan an executable function associated with a first operation code; and

(b) augmenting said first operation code in the access plan with a pointer to said executable function.

2. The method as claimed in claim 1, further comprising repeating steps (a) and (b) for the remaining operation codes in the access plan.

3. The method as claimed in claim 1, wherein said step (b) comprises augmenting said first operation code in the access plan with a pointer to an intermediate function, said intermediate function including a data structure for storing a pointer to said executable function.

4. The method as claimed in claim 3, wherein said data structure includes means for storing information associated with said executable function or said first operation code.

5. The method as claimed in claim 1, wherein said step (b) comprises augmenting said first operation code in the access plan with a second pointer to a data structure, said data structure providing means for storing information associated with said first operation code or said executable function.

6. The method as claimed in claim 1, wherein said step (a) further includes assessing the

executable function associated with the first operation code and if applicable, replacing the call to the executable function with a call to a second executable function.

7. The method as claimed in claim 3, wherein said intermediate function includes processing operations for the first operation code or the executable function associated with the first operation code.

8. The method as claimed in claim 7, wherein said processing operations in the intermediate function include gathering statistics on the use of the executable function associated with the operation code.

9. The method as claimed in claim 7, wherein said processing operations in the intermediate function include a pause for receiving user input before or after the call to the executable function.

10. A computer program product for use on a computer wherein queries are entered by a user for retrieving data in a relational database management system having a query optimizer for generating an access plan for executing the query, said computer program product comprising:

a recording medium;

means recorded on said recording medium for instructing said computer to perform the steps of:

(a) determining an executable function associated with a first operation code in the access plan, the first operation code being one of a plurality of operation codes; and

9 (b) augmenting said first operation code in the access plan with a pointer to said  
10 executable function.

1 11. The computer program product as claimed in claim 10, the means for instructing said  
2 computer further comprising repeating steps (a) and (b) for the remaining operation codes in the  
3 access plan.

1 12. The computer program product as claimed in claim 10, wherein said step (b) comprises  
2 augmenting said first operation code in the access plan with a pointer to an intermediate function,  
3 said intermediate function including a data structure for storing a pointer to said executable  
4 function.

1 13. The computer program product as claimed in claim 12, wherein said data structure  
2 includes means for storing information associated with said executable function or said first  
3 operation code.

1 14. The computer program product as claimed in claim 10, wherein said step (b) comprises  
2 augmenting said first operation code in the access plan with another pointer to a data structure,  
3 said data structure providing means for storing information associated with said first operation  
4 code or said executable function.

1 15. The computer program product as claimed in claim 10, wherein said step (a) further  
2 includes assessing the executable function associated with the first operation code and if  
3 applicable, replacing a call to the executable function with a call to another executable function.

16. The computer program product as claimed in claim 12, wherein said intermediate function includes processing operations for the first operation code or the executable function associated with the first operation code.

17. The computer program product as claimed in claim 16, wherein said processing operations in the intermediate function include gathering statistics on the use of the executable function associated with the first operation code.

18. The computer program product as claimed in claim 12, wherein said processing operations in the intermediate function include a pause for receiving user input before or after a call to the executable function.

19. A relational database management system for use with a computer system wherein queries are entered by a user for retrieving data from tables, the relational database management system including a query optimizer for generating an access plan associated with the queries entered by the user, said relational database management system comprising:

(a) means for determining an executable function associated with each of a plurality of operation codes in the access plan; and

(b) means for augmenting said operation codes in the access plan with a pointer to said executable function associated with each operation code.

20. The relational database management system as claimed in claim 19, wherein said means for augmenting said operation codes includes means for replacing said operation codes in the access plan with a pointer to an intermediate function, said intermediate function including a data

4 structure for storing a pointer to said executable function.

1 21. The relational database management system as claimed in claim 20, wherein said data  
2 structure includes means for storing information associated with said executable function or said  
3 operation codes.

1 22. The relational database management system as claimed in claim 19, wherein said means  
2 for augmenting said operation codes includes means for adding another pointer to a data  
3 structure, said data structure providing means for storing information associated with said  
4 operation codes or said executable function.

ABSTRACT

DIRECT CALL THREADED CODE

A pre-pass and direct call mechanism which replaces the interpreter. The mechanism moves static decisions out of the repeated runtime path into a pre-pass operation.

- 5 Advantageously, the mechanism reduces runtime overhead and improves overall performance of the DBMS during runtime, while maintaining the considerable investment in opcode generation and plan representation.






FIG. 1

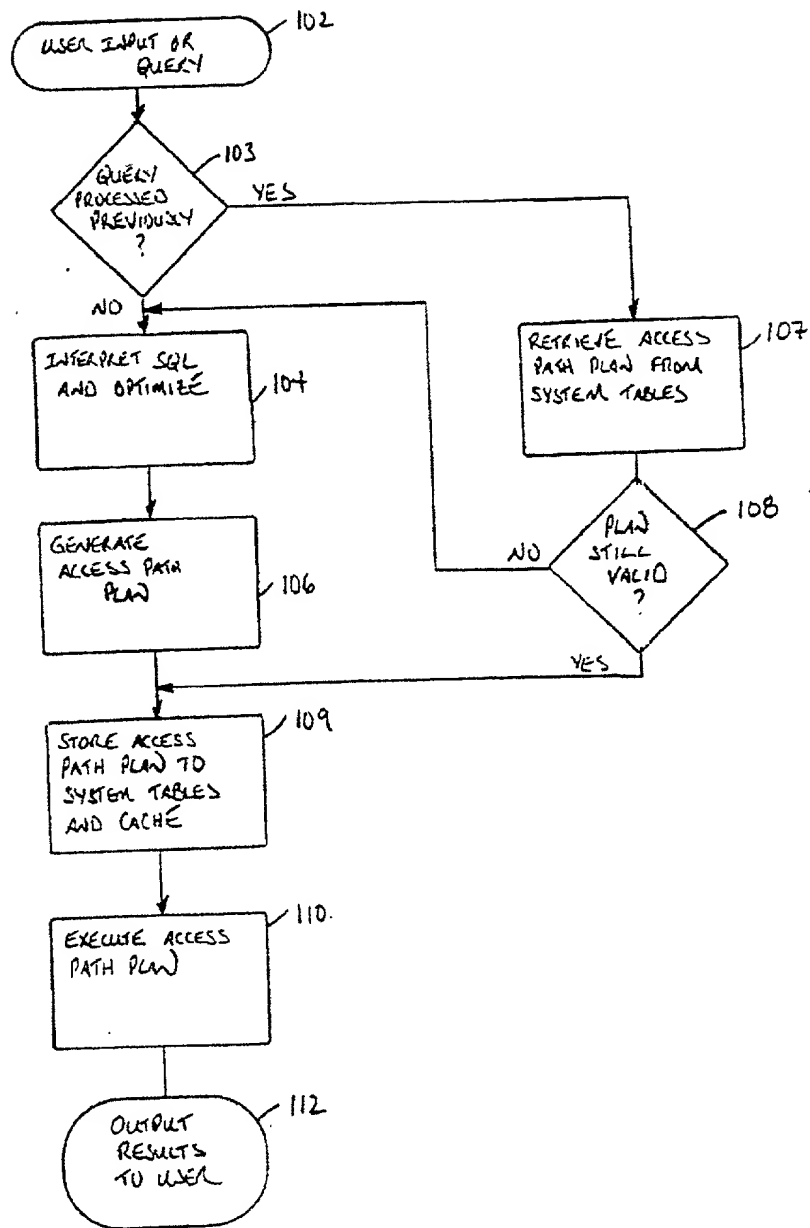
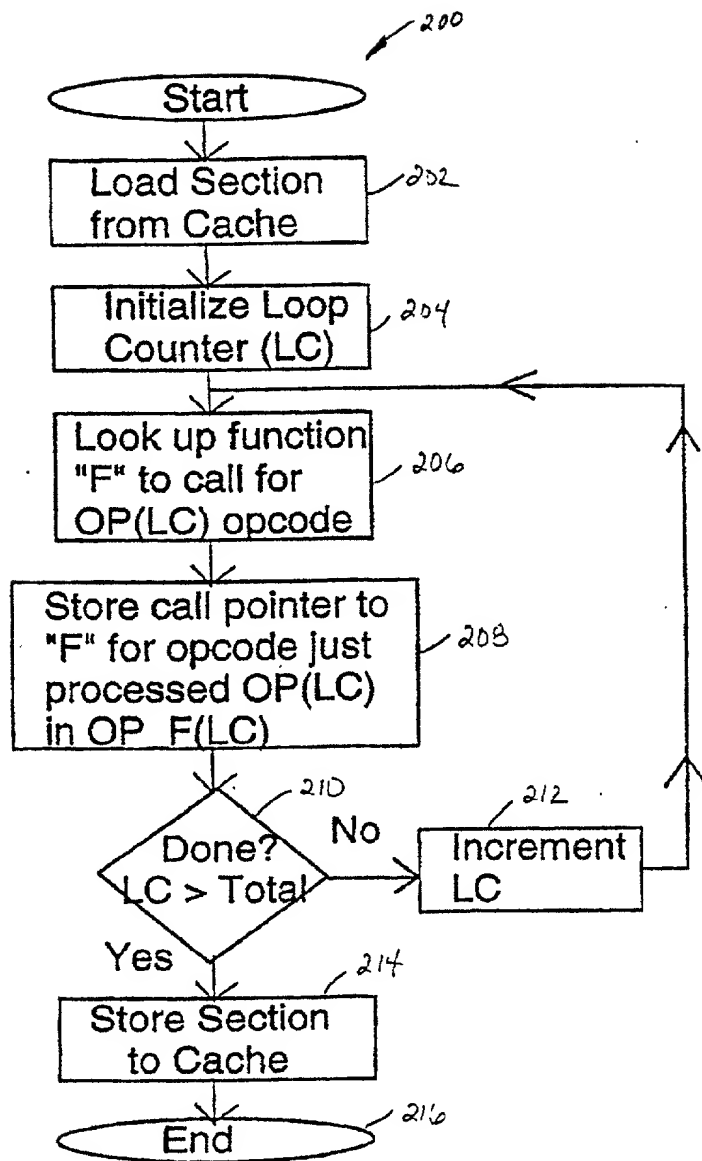
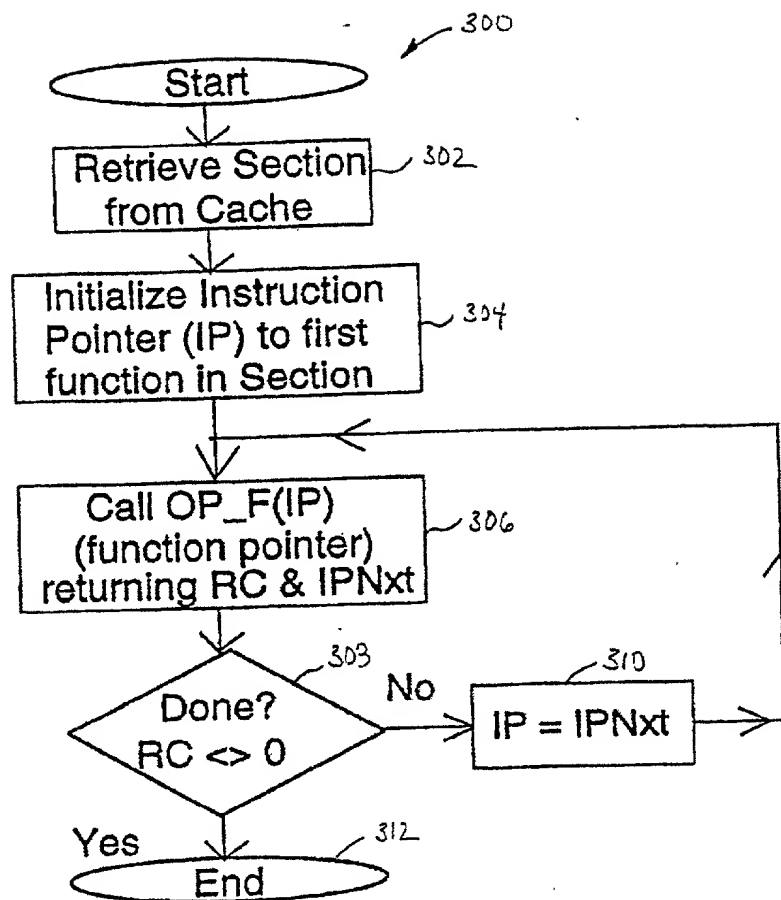


FIG. 2



Translation of Opcode to Function Pointer

FIG. 3



New implementation of Section Execution

FIG. 4

**DECLARATION, POWER OF ATTORNEY AND PETITION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled DIRECT CALL THREADED CODE the specification of which

(check one) ☒ is attached hereto.  
☐ was filed on \_\_\_\_\_ as Application Serial No. \_\_\_\_\_  
 and was amended on \_\_\_\_\_ (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, § 1.56(a).

I hereby claim foreign priority benefits under Title 35, United States Code, § 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application of which priority is claimed.

**PRIOR FOREIGN APPLICATION(S)**

	Priority Claimed	
2,279,222; Canada; July 30, 1999	<u>X</u>	<u>      </u>
(Number)(Country)(Day, month, year filed)	Yes	No
_____	<u>      </u>	<u>      </u>
(Number)(Country)(Day, month, year filed)	Yes	No
_____	<u>      </u>	<u>      </u>
(Number)(Country)(Day, month, year filed)	Yes	No

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, § 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial No.)	Filing Date	(Status: Patented, pending, abandoned)
_____	_____	_____
(Application Serial No.)	Filing Date	(Status: Patented, pending, abandoned)
_____	_____	_____

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

And I hereby appoint W. POMS, Reg. 18,782; G.P. SMITH, Reg. 20,142; A.C. ROSE, Reg. 17,047; L.J. BOVASSO, Reg. 24,075; D.J. OLDENKAMP, Reg. 29,421; C. DARROW, Reg. 30,166; M.E. HARRIS, Reg. 26,690; K.A. MACLEAN, Reg. 31,118; C. ROSENBERG, Reg. 31,464; M.E. BROWN, Reg. 28,590; A. KONDZELLA, Reg. 18,013; A.P. BLOCK, Reg. 35,450; S.R. HANSEN, Reg. 38,486; D.N. LARSON, Reg. 29,401; J.W. INSKEEP, Reg. 33,910; G.F. COTTERELL, Reg. 37,802; H.D. JASTRAM, Reg. 19,777; J.M. KALIS, Reg. 37,650; C.J. LERVICK, Reg. 35,244; C.A.S. HAMRICK, Reg. 22,586; R.O. GUILLOT, Reg. 28,852; C. CHOU, Reg. 41,672; A. DIEPENBROCK, Reg. 39,960; M. BOSWORTH, Reg. 28,186; R. STRIMAITIS, Reg. 35,697; P.W. JOHNSON, Reg. 33,123; I.M. FOERSTER, Reg. 36,511; T.M. FARRELL, Reg. 37,321; C.A. HUGHES, Reg. 26,914; J.E.H. HOEL, Reg. 26,279; E.A.P. PENNINGTON, Reg. 32,588; J.C. REDMOND, Reg.

18,753; J.F. VILLELLA, JR., Reg. 30,599; R.J. BLUESTONE, Reg. 40,518; C. BERMAN, Reg. 29,249; G.B. WOOD, Reg. 29,133; M.B. FARBER, Reg. 32,612 and S.J. LAZARIS, Reg. 45,981, OPPENHEIMER WOLFF & DONNELLY LLP, 101 Park Center Plaza, Suite 400, San Jose, California 95113, (408) 795-3000, as my attorneys with full power of substitution and revocation, to prosecute said application and to transact in connection therewith all business in the Patent and Trademark Office and before competent International Authorities.

Address all telephone calls to Spyros J. Lazaris at (310) 788-5000, and address all correspondence to:

SPYROS J. LAZARIS, ESQ.  
**OPPENHEIMER WOLFF & DONNELLY LLP**  
2029 CENTURY PARK EAST, 38TH FLOOR  
LOS ANGELES, CA 90067

Wherefore I pray that Letters Patent be granted to me for the invention or discovery described and claimed in the foregoing specification and claims, and I hereby subscribe my name to the foregoing specification and claims, declaration, power of attorney, and this petition.

Full Name of Sole or First Inventor: Ian R. Finlay

Home Address: 514 Davis Drive, R.R. #1, Uxbridge, Ontario, Canada L9P 1R1

Post Office Address: Same as above

Citizenship: Canada

Inventor's Signature: \_\_\_\_\_ Date: \_\_\_\_\_

Full Name of Second Inventor: Douglas James Doole

Home Address: #1718, 45 Wynford Heights Cresc., Don Mills, Ontario, Canada M3C 1L3

Post Office Address: Same as above

Citizenship: Canada

Inventor's Signature: \_\_\_\_\_ Date: \_\_\_\_\_